ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THPV049

VIRTUALISATION AND SOFTWARE APPLIANCES AS MEANS FOR DEPLOYMENT OF SCADA IN ISOLATED SYSTEMS

P. Golonka†, L. Davoine, M. Zimny, L. Zwalinski, CERN, Geneva, Switzerland

Abstract

The paper discusses the use of virtualisation as a way to deliver a complete pre-configured SCADA (Supervisory Control And Data Acquisition) application as a software appliance to ease its deployment and maintenance. For the off-premise control systems, it allows for deployment to be performed by the local IT servicing teams with no particular control-specific knowledge, providing a "turnkey" solution. The virtualisation of a complete desktop allows to deliver and reuse the existing feature-rich Human-Machine Interface experience for local operation; it also resolves the issues of hardware and software compatibilities in the deployment sites. The approach presented here was employed to provide replicas of the "LUCASZ" cooling system to collaborating laboratories, where the on-site knowledge of underlying technologies was not available and required to encapsulate the controls asa "black-box" so that for users, the system is operational soon after power is applied. The approach is generally applicable for international collaborations where control systems are contributed and need to be maintained by remote teams.

MOTIVATION

For the past two decades industrial controls technologies have been applied with success at CERN to build control systems for a vast range of applications, ranging from laboratory setups through technical infrastructure up to the ones for accelerators and large particle detectors counting millions of I/O channels. Applying the standardized stack of industrial technologies (WinCCOA SCADA [1]), enhanced with in-house developed frameworks [2,3] allowed for rapid development and effective maintenance of hundreds of complex control system across the organisation. However, this wouldn't be possible without centrally supported IT services and infrastructure. In particular, provisioning of secure network connectivity, computing server hardware, installation of operating system and software, high capacity database service, shared file-systems, consoles in the control rooms or terminal servers for secure remote access have been essential for reliable day-to-day operation. Harmonized efforts of numerous expert teams, ensured the maintenance and upgrades of applications and infrastructure [4] throughout the life-time of control systems.

In general, production control systems are hosted on centrally maintained infrastructure, remain highly homogeneous in terms of technologies and assume the availability of services and also experts. Nevertheless, this may not be achievable for international research collaborations, where the development and precommissioning of the control system may be performed in collaborations, where the development collaborations, where the development and precommissioning of the control system may be performed in a different place than the final location of the equipment. The local infrastructure, available expertise and policies at the deployment site may differ, making the integration and maintenance of such control system seemingly impossible.

LUCASZ CO2 Cooling Stations

LUCASZ [5], *Light Use Cooling Appliance for Surface Zones*, is a medium size I-2PACL CO_2 cooling system developed at CERN with 1kW cooling power at −35°C. It was designed with the purpose of cooling down small and medium sized light tracking detector structures and supporting detector-assembly. A number of research institutes collaborating with CERN needs to build their own replicas of the LUCASZ system to proceed with their activities and commitments related to the phase-II upgrade of the LHC Experiments.

The control system for LUCASZ employs the standard CERN UNICOS [3] technology stack with a Schneider M580 PLC. To facilitate the local operation for non-expert users LUCASZ has a detachable *LocalBOX* featuring an industrial touch-panel. However, its functionalities are limited: for instance it is not suitable for long-term storage of historical data, hence it may not be considered a replacement for a complete, PC-based UNICOS SCADA.

Provisioning a stand-alone cooling system to be operated outside CERN is a technical challenge and requires expertise in several domains not always available in the institutes. Indeed, in addition to cooling, mechanical engineering also the expertise in deployment and configuration of control systems would be required.

PROBLEM AND PROPOSED SOLUTION

Difficulties in deployment and maintenance of control systems, requiring specialized skills, may hinder the effective collaboration and prevent possible re-use of existing instruments or technologies. A solution allowing for the control system to be deployed in form of a "black box" next to the equipment, with the purpose of acting as the *local control station* is demanded.

Unlike for the PLCs, deploying the full UNICOS SCADA application on top of undefined IT infrastructure remains a complex task requiring specific knowledge. In what follows we propose a solution allowing to package a complex SCADA application together with the necessary environment, and at the same time address the problem of the infrastructure compatibility.

Virtual Software Appliances

To address the class of practical problems of lack of local domain-specific technical expertise the IT industry applies a pattern of *software appliances* [6]: self-contained

† Piotr.Golonka@cern.ch

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THPV049

assemblies of software tailored and configured for specific use and easy to deploy through well known and practised IT procedures.

Virtual machines (*VM*s) are a widely used method to implement software appliances, supported by IT industry. By virtues of *isolation* and *emulation* they allow to decouple the physical computing infrastructure from the one required by the software embedded in the appliance.

Instances of VMs are created from *VM image files* comprising all the software, the operating system and settings. The image files are the standardized [7] way of distributing the appliances from vendors to end-users.

Local Desktop Virtualisation

Addressing the primary requirement of *local* operation of SCADA, the proposed solution adopts the approach of *local desktop virtualisation* to provide a complete desktop session to the operators using the VM running in a local desktop PC located close to the controlled equipment (see Fig. 1).

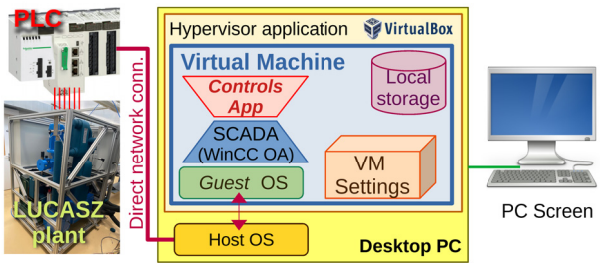

Figure 1: The architecture of proposed solution that employs local desktop virtualisation.

This approach draws from a decade of authors' experience with virtualisation of Linux-based desktop developer environments using the multi-platform and open-source *VirtualBox* hypervisor [8]. Image files containing pre-configured tools and services allow to quickly instantiate new and ready-to-use VMs to be used by developers, in particular those with no experience with configuring a Linux desktop machine. It enables them to develop and test the Human-Machine Interface (HMI) with the user experience similar to real operator consoles, considering aspects such as colour themes, available fonts or pixel-exact rendering. The VMs behave identically regardless of the operating system (Windows, MacOS, various Linux distributions) installed on their desktop machine or a laptop (hence prove the capability of this technology to address the requirement of decoupling the appliance from the hosting infrastructure).

In practice, the perceived performance of VMs hosted in modern desktop machines equipped with SSD storage and enough RAM (extra 4 GB per VM) is most often indistinguishable from the non-virtualised environments. However it is essential to activate the acceleration of virtualisation (VT-d, IOMMU) in the computer's setup – options which are usually disabled by default.

The VM images are based on the currently recommended version of the operating system and applications (CentOS 7, WinCC OA 3.16). New releases containing updates/patches applied are delivered on a regular basis and their import takes around 2 minutes. Desktop session (KDE Plasma) is started immediately after the machine boots, ready to be used in less than 30 seconds. The users appreciate simplicity of use and additional features: auto-adjusting the VM's screen resolution upon its window resize, clipboard synchronization, access to host files, and adequate functionality even without a network connection.

Using these already existing images as a starting point we were able to prototype a "black-box" SCADA appliance for LUCASZ with not much effort. The initial experience clearly indicated that the approach would fit the needs remarkably well, while reusing the already existing work.

Approach to Computer Security

Through the *isolation* features, virtualisation may significantly enhance the security of deployed appliances, if configured correctly. In what follows we assume that the security of the host machine: not only the OS but also the hypervisor software (e.g. VirtualBox) is assured adequately (security and functionality erratas, etc.).

To reduce the potential surface for vulnerabilities the amount of installed software and enabled services in the appliance is reduced to strict minimum. This also allows to reduce the VM startup time and the size of the distributed image (2.5 GB).

The appliance does not require any external *online* resources and could be operated in complete isolation.

A direct connection to the PLC using a dedicated Ethernet network port of the host PC is required. The VM configures it in *bridged mode* for its own use and it should not be available for any other purpose by the host. The MAC address of the network port is pre-configured by the VM, and static non-routable IP addresses are configured for the PLC and the VM's network stack.

If necessary, inter/intra-net connectivity from VM may be enabled: a separate network connection is preconfigured in the VM to provide routed networking through the host's networking stack working in the NAT (*Network Address Translation*) mode. When enabled, it provides the security similar to those of a home router, disallowing incoming connection to the VM, and having all outgoing connection handled by the host's internet connection, respecting all security policies configured by its administrators.

The operation of the appliance is primarily through the main window of the VirtualBox application, which displays the contents of virtualised screen, and passes the mouse and keyboard events to the VM. Remote assistance or operation to the session in which the VM executed may be provided either by the standard tooling or via a commercial *VirtualBox Extension* providing direct RDP access to its console.

The appliance is ready for operation within 30 seconds from its start. Once the VM loads the operating system it auto-starts the SCADA and opens the main HMI Window allowing to visualize the state of the control system. The SCADA-level access control mechanisms are employed requiring the operator to log in to enable commanding on the controlled hardware.

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THPV049

Storage for Historical Values Data

To store the history of the value-changes the SCADA in the self-contained appliance may not assume the availability of large centrally-managed databases, as these typically used at CERN, and needs to fall back to the local file-based mechanism available in WinCC OA. This requires careful parametrisation of data-retention and data-reduction policies to balance the period of available data, its granularity and consumed storage space.

In the proposed solution a dedicated disk image file, with (self-expanding) capacity of up to 100GB is attached to the VM and auto-configures itself to become the storage space for the archive data. The image may be freely detached from one VM and re-attached to another allowing for upgrades of the VM without the loss of archives. In addition, detaching it prior to a VM snapshot allows to exclude the historical data and reduce its size.

Reliability and Recovery Procedure

The reliability of a VM depends of the hardware it runs on. For stand-alone control systems it is not optimal to deploy high-end server hardware. Often a cheaper PC may be sufficient, even though their reliability is lower, provided that a quick and straightforward recovery procedure is available. In case one needs to change the PC, it is sufficient to load the same VM *image* used for the original installation, which takes a few tens of seconds. It is also possible to re-connect the storage of historical values. In addition, standard functionality of VM *snapshots*, which themselves could be exported as images, provide a robust solution for backups.

Update Scenarios

In spite of being a "black box" solution, the appliance needs to be managed through the lifetime of the LUCASZ system. To assure proper support, the major versions of the SCADA and OS of the appliance need to be kept in line with those used at CERN. The envisaged scenario is to prepare and deliver new versions of the image on, with updated components on a regular basis (once per year). The new image may be imported as a new VM working *alongside* the one used in production allowing for its local commissioning and smooth roll-back scenarios. At the moment of writing we have no experience yet with executing an upgrade of the image in production.

Discussion of Alternatives

The presented choice of approach and technologies may not be the only one allowing to address the use case. In what follows, let us present the reasons for not having followed some of other possible solutions.

Remote access to appliances hosted in the on-premises or *cloud* data centres is the most widely known provisioning method for desktop virtualisation. At CERN it is used to provide remote operation of the applications using Terminal Servers hosted on the local "*openstack*" cloud infrastructure.

However, the essential part for the cloud models (Infrastructure/Platform/Software/Desktop as a Service) is the availability of network communication between the user and the hosting site. Concerns about the privacy and non-locality of data may arise due to internal policies. The requirement of vicinity (direct connection) of the self-contained operator station to the operated plant may be hard to address. In addition, the efforts and skills related to provisioning and configuration of such services for a single appliance may exceed those needed by the presented solution.

IIoT and *edge* solutions are very popular trends in industrial controls these days. However, they are not applicableto the discussed problem, as there is no need for storing and processing data in a central place anyway.

Containers are amongst the most popular methods of deployment for reusable software appliances, also applied for Linux-based control systems at CERN [9]. The upcoming enhancements in MS Windows container subsystem allowing to run Linux containers with GUI may become a viable alternative, with advantage of significant reduction in the size of distributed image and startup time, yet at the cost of reduced level of isolation.

Centralized deployment and maintenance of controls applications are being addressed by the SCADA vendors too. For WinCC OA the "*On Premise Administration*" feature (part of *IOT Suite* [10]) has recently become available, yet at the moment of this writing we have no practical experience with it yet.

The fall-back solution is the provisioning of the physical PC, configured with all the software to be installed in the remote location. However, remote expert interventions may be restricted by local IT policies.

Numerous available software products allow for effective virtualisation of GUI desktops in local machines, with efforts for standardization and interoperability [7]. Our choice of Virtual Box was motivated by its availability on numerous platforms and being an open source project. We are convinced that in a general case, other products could fit the same purpose.

EXPERIENCE OF FIRST DEPLOYMENT

The deployment of a virtualised SCADA has been successfully completed for the first time with the LUCASZ installation at DESY in Hamburg, Germany (Fig. 2).

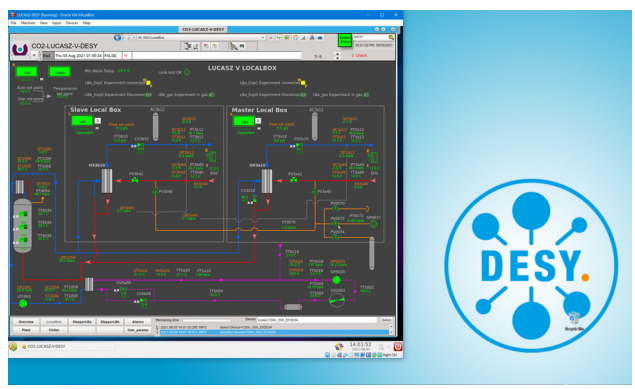

Figure 2: Operator screen of the *CO2-LUCASZ-V-DESY* appliance visible in the VM window of a desktop PC.

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THPV049

Prior to the deployment, the image of the appliance has been prepared by the experts at CERN. Firstly, a *multipurpose image* with latest versions of SCADA software, OS patches, and a complete generic unparameterized UNICOS application and preconfigured networking parameters (for the PLC communication) was prepared. Based on that the LUCASZ control engineers instantiated a VM, configured the SCADA application (import device instances, I/O addresses, application-specific customizations and operator panels), validated it with a PLC and hardware. Then the VM was exported as the *final image* and delivered to the IT team in the remote institute to be instantiated in the local PC. The same image may be used (with minimal adjustments) for all other instances of the LUCASZ cooling plant.

The initial deployment at DESY resulted in successful operation: after loading and connecting the PLC (with CERN remote assistance), the communication with SCADA was established immediately. The deployment and hand-over to the local team for operation was completed in less than a day, which is very encouraging for the future.

APPLICABILITY

The presented solution addresses primarily the need for an easy-to-deploy standalone local control system to be used for replicas of highly specialized equipment, built from engineering *blueprints.* However, the scope of the use for this approach is wider.

Large and complex research equipment is often built through a collaborative effort partitioned to work-packages distributed worldwide. Some of the contributions may be delivered as complete "turn key" systems to bring also the *operational experience* (e.g. gas system for KEK contributed by CERN). Their controls technologies may differ from the ones used widely by the collaboration whereas re-engineering is not feasible. Virtualised appliances allow to overcome this inhomogeneity and ease the integration.

Recently the appliance has been used to implement a small portable laboratory setup at CERN for tests of CO_2 long vertical transfer lines recording data from a few pressure sensors wired to a PLC. It took less than 3 hours to set up the PLC and the SCADA appliance running in a laptop. Numerous CERN groups have already expressed their interest for this type of portable SCADA deployments to be used in instrumentation of test stands.

CONCLUSIONS AND OUTLOOK

The use of virtualisation to provision self contained SCADA systems in form of easy-to-deploy software appli-ances has been carried out successfully for the first remote instance of LUCASZ system in DESY, with subsequent ones already scheduled for NIKHEF and INFN. The same approach is envisaged to provide control system for the Gas System to be contributed by CERN to the T2K-ND280 experiment in Japan. After a decade of the internal use of desktop virtualisation for software development, we found the approach to be effective for production and lab setups to deliver the full functionality of SCADA. We expect it to be used more and more not only internally at CERN but also for contributions to international research collaborations.

We believe that the presented solution is generally applicable also for other control system technologies and will further enable the reuse of unique hardware across research community, also beyond physics.

REFERENCES

[1] Siemens Simatic WinCC Open Architecture, www.winccoa.com

[2] O. Holme, M. Gonzalez Berges, P. Golonka, S. Schmeling, "The JCOP framework", in *Proc. ICALEPCS'05*, Geneva, Switzerland, Oct. 2005, paper WE2.1-6O.

[3] E. Blanco, F. B. Bernard, P. Gayet, H. Milcent, "UNICOS: An Open Framework", in *Proc. ICALEPCS'09*, Kobe, Japan, Oct. 2009, paper THD003, pp. 910-912.

[4] R. Kulaga, J. Arroyo Garcia, M. Boccioli, E. Genuardi, and P. Golonka, "Large-scale upgrade campaigns of SCADA systems at CERN – organisation, tools and lessons learned" in *Proc. ICALEPCS'17*, Barcelona, Spain, Oct 2017, pp. 1384-1388. doi:10.18429/JACoW-ICALEPCS2017-THPHA021

[5] L. Zwalinski *et al.*, "First steps in automated software development approach for LHC Phase II upgrades CO_2 detector cooling systems", in *Proc. ICALEPCS'19* , New York, USA, Oct 2019, pp. 1488-1491. doi:10.18429/JACoW-ICALEPCS2019-WEPHA170

[6] "Virtual Appliances: A New Paradigm for Software Delivery", VMWare whitepaper, https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vam/vmware-virtual-appliance-solutions-white-paper.pdf

[7] Open Virtualization Format, DTMF specification, ISO-17203, https://www.dmtf.org/standards/ovf

[8] VirtualBox cross-platform virtualisation application, www.virtualbox.org

[9] R. Voirin, T. Oulevey, and M. Vanden Eynden, "The state of containerization in CERN Accelerator Controls", presented at ICALEPCS'21, Shanghai, China, October 2021, paper THBL03, this conference.

[10] Siemens Simatic WinCC OA IOT Suite, https://new.siemens.com/global/en/products/automation/industry-software/automation-software/scada/simatic-wincc-oa/simatic-wincc-oa-iot-suite.html